arXiv:1510.03300v1 [physics.med-ph] 9 Oct 2015

The scenario-based generalization of radiation therapy margins

Albin FREDRIKSSON*† and Rasmus BOKRANTZ†

October 13, 2015

Abstract

We give a scenario-based treatment plan optimization formulation that is equivalent to planning with geometric margins if the scenario doses are calculated using the static dose cloud approximation. If the scenario doses are instead calculated more accurately, then our formulation provides a novel robust planning method that overcomes many of the difficulties associated with previous scenario-based robust planning methods. In particular, our method protects only against uncertainties that can occur in practice, it gives a sharp dose fall-off outside high dose regions, and it avoids underdosage of the target in "easy" scenarios. The method shares the benefits of the previous scenario-based robust planning methods over geometric margins for applications where the static dose cloud approximation is inaccurate, such as irradiation with few fields and irradiation with ion beams. These properties are demonstrated on a suite of phantom cases planned for treatment with scanned proton beams subject to systematic setup uncertainty.

1. Introduction

Safety margins are the most utilized means of generating radiation therapy treatment plans that are robust against errors. For instance, to mitigate the risk that the target is located differently during the plan delivery than during the image acquisition, a planning target volume (PTV) is usually defined as a geometric expansion of the clinical target volume (CTV), and planning is performed to deliver high dose to the full PTV. Analogously, the uncertainty in the position of an organ at risk (OAR) can be mitigated by an expansion into a planning organ-at-risk volume (PRV). Recipes for the calculation of margins on the basis of systematic and random uncertainties and the amplitude of respiration are surveyed in van Herk [21].

An underlying assumption when geometric margins are used is that the static dose cloud approximation is accurate. According to this approximation, the patient is moving (as an effect of the errors) inside a static dose distribution that is not affected by the motion [19]. This approximation is often sufficiently accurate for tumors in regions of relatively homogeneous density, such as cancers in the pelvic region [4]. The approximation is, however, inaccurate for sites of heterogeneous density, such as lung, which can lead to under- or overdosage if the colocation of the target and the treatment fields differs from what was planned [4]. The approximation is also inaccurate for treatments with few fields, because it does not model that a setup shift in parallel with a beam axis only affects the corresponding beam dose marginally

*E-mail: albin.fredriksson@raysearchlabs.com
†RaySearch Laboratories, Sveavägen 44, SE-111 34 Stockholm, Sweden.

(mainly by scaling according to the inverse-square law), whereas perpendicular shifts affect the beam dose to a large extent (because the dose is displaced). The situation is aggravated for ion treatments, the depth-dose curves of which are more sensitive to the density distribution of the treated volume than those of photon treatments [13].

As alternatives to geometric margins, methods that take uncertainties into account explicitly in the optimization have been proposed [8, 16, 19, 20]. In these methods, dose distributions for multiple error scenarios (e.g., setup shifts, scaled densities, or displaced organs) are computed and the treatment plans are optimized with respect to all of these doses simultaneously. The methods differ in how they take the dose distributions into account: expected value optimization [19, 20] minimizes the expectation of the objective value over the scenario doses, composite worst case optimization [8] minimizes the objective value of the worst case scenario, and voxelwise worst case optimization [16, 20] minimizes the objective function applied to a worst case dose distribution defined as the worst scenario dose to each voxel considered independently. While these methods can produce robust plans also for sites of heterogeneous density, they result in qualitatively different plans than margin-based planning, even in situations where the static dose cloud approximation holds. The differences are not always desirable [6, 7]: expected value optimization protecting against systematic errors results in gradual dose fall-offs, something that only marginally increases the probability of tumor control but still results in highly increased integral dose; composite worst case optimization risks neglecting "easy" scenarios; and voxelwise worst case optimization does not utilize that an increased dose and the periphery of the high-dose region leads to a steeper dose gradient against the non-involved tissue and is overly conservative when used in combination with dose-volume histogram (DVH) goals, because it considers an unphysical dose distribution.

In the present paper, we introduce a generalization of margin-based planning. The generalization is scenario-based, but does not suffer from the disadvantages mentioned above. When the static dose cloud approximation is exact, our method is equivalent to planning with geometric margins. However, the method can also be used to achieve robust plans when the static dose cloud approximation is inaccurate. Because the static dose cloud approximation is implicitly assumed with the use of geometric margins, this implies that our method is a generalization of treatment plan optimization with margins.

2. Methods

We generalize the concept of treatment plan optimization with margins to the case when scenario doses are computed for multiple scenarios. The following exposition concerns a single CTV subject to setup errors, but the method can be readily applied to multiple regions of interest (ROIs), including OARs. For simplicity of the presentation, it is assumed that all setup error scenarios correspond to geometric shifts that can be represented by vectors with all components being multiples of the voxel dimensions, and that the voxels of the considered ROI lines up with the dose grid.

The exposition is structured as follows: We first highlight that the use of margins entails the implicit assumption that the static dose cloud approximation holds. We then show how a voxel-separable plan evaluation criterion applied to a PTV can be reformulated as a criterion applied to the CTV under multiple scenarios with scenario doses given by the static dose cloud approximation. Using this reformulation, we define a scenario-based criterion that utilizes the true, non-approximate, scenario doses—this is the scenario-based generalization of margins. Finally, we give a number of possible refinements of this criterion and discuss how to adapt the method to other criteria than directly voxel-separable ones.

2.1. Margins implicitly assume the static dose cloud approximation

Target margins are generally defined using a geometric expansion of the CTV, and planning is performed to deliver high dose to the PTV, which is the union of the CTV and the margin [10, Section 4.5]. The size of the PTV is selected so that the probability that the CTV will be within the PTV is high [23]. For this to imply a high probability of the CTV receiving a high dose—which is the ultimate goal of the use of a margin—the static dose cloud approximation must be reasonably accurate: If it is, then whenever the CTV displaced by the errors is geometrically within the PTV, it will receive a high dose. If the approximation on the other hand is inaccurate, meaning that the dose might deform as an effect of the displacement, then there are no guarantees regarding the dose received by the moving CTV. This is the reason behind the failure of geometric margins in highly modulated proton therapy [8, 13].

2.2. A scenario-based reformulation of margins

We consider an arbitrary plan evaluation criterion that is additively separable over the voxels. Let the voxelwise component of the criterion be given by the function ϕ. Applied to the PTV, the criterion takes the form

$$\sum_{i \in \mathcal{P}} \phi(d_i), \tag{2.1}$$

where d_i is the dose to voxel i and $\mathcal{P}$ an index set over the PTV voxels. Common voxel-separable criteria include minimum, maximum, and uniform dose penalties, see, e.g., Oelfke and Bortfeld [15].

According to the static dose cloud approximation, the effect of a setup error is that the patient is shifted rigidly inside a static dose distribution [19]. A rigid shift of the patient corresponds to the voxel indices being shifted by a fixed offset (assuming that the shift lines up with the dose grid). Let $\mathcal{S}$ denote the set of scenarios and let the offset corresponding to the shift s, for s in $\mathcal{S}$, be denoted by $j(s)$. The CTV voxels under scenario s are then given by the scenario-dependent set $\mathcal{C}(s) := \{i + j(s) : i \in \mathcal{C}\}$, where $\mathcal{C}$ enumerates the CTV voxels in the nominal scenario (corresponding to no error). Because the PTV is an expansion of the CTV, it holds that $\mathcal{P} = \cup_{s \in \mathcal{S}} \mathcal{C}(s)$ if $\mathcal{S}$ is appropriately defined, as illustrated in Figure 1. The criterion (2.1) can thus be equivalently stated as

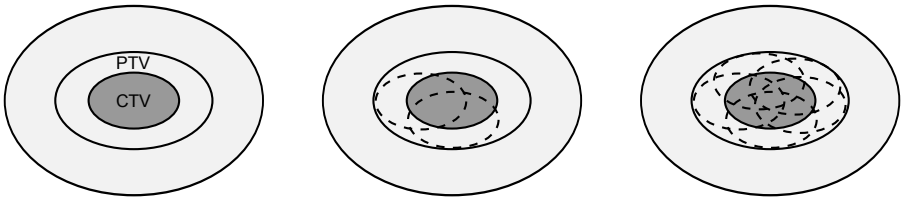

(a) A patient with a CTV and a PTV.

(b) The dashed ellipses illustrate $\mathcal{C}(s) = \{i + j(s) : i \in \mathcal{C}\}$ for two scenarios s.

(c) The set $\cup_{s\in\mathcal{S}}\mathcal{C}(s)$ approaches $\mathcal{P}$ as the number of scenarios increases.

Figure 1. Construction of the PTV from translated copies of the CTV.

$$\sum_{i\in\bigcup_{s\in\mathcal{S}}\mathcal{C}(s)} \phi(d_i).$$

This expression, in turn, is equivalent to

$$\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{C}(s)} p_i\phi(d_i),$$

where p_i is the reciprocal of the number of scenarios s under which voxel i is in $\mathcal{C}(s)$:

$$p_i = 1/\left|\{s \in \mathcal{S} : i \in \mathcal{C}(s)\}\right|,$$

where $|\cdot|$ denotes the cardinality of a set, see Figure 2.

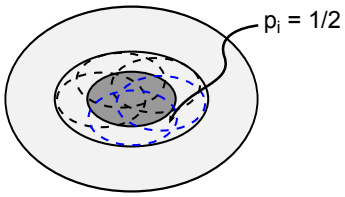

Figure 2. Voxel i has weight $p_i = 1/2$ because it is contained in $\mathcal{C}(s)$ (dashed ellipses) under two scenarios s.

Offsetting the voxel indices within a fixed dose distribution is equivalent to keeping the voxel indices fixed but shifting the dose rigidly. The criterion can therefore also be formulated as

$$\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{C}} p_{i+j(s)}\phi(d_{i+j(s)}).$$

Let $\tilde{d}_i(s) := d_{i+j(s)}$ denote the dose to voxel i under scenario s and $p_{i,s} := p_{i+j(s)}$ denote the corresponding weight. The criterion can then be written as

$$\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{C}} p_{i,s}\phi(\tilde{d}_i(s)), \tag{2.2}$$

which is our scenario-based reformulation of the PTV criterion.

In summary, we have derived the following equivalence:

$$\sum_{i\in\mathcal{P}} \phi(d_i) = \sum_{s\in\mathcal{S}} \sum_{i\in\mathcal{C}} p_{i,s}\phi(\tilde{d}_i(s)), \tag{2.3}$$

where $\tilde{d}_i(s) = d_{i+j(s)}$ and

$$p_{i,s} = 1/\left|\left\{s' \in \mathcal{S} : i + j(s) - j(s') \in \mathcal{C}\right\}\right|, \tag{2.4}$$

and $\mathcal{S}$ has been selected such that $\mathcal{P} = \{i + j(s) : i \in \mathcal{C}, s \in \mathcal{S}\}$.

2.3. The scenario-based generalization of margins

The expression (2.2) provides a reformulation of the criterion (2.1) for the PTV to a scenario-based criterion for the CTV. The static dose cloud approximation was used to define the weights $p_{i,s}$ and the dose distribution $\tilde{d}(s)$. We propose the use the scenario-based criterion (2.2) together with the exact dose distribution (or a more accurate approximation than the static dose cloud approximation) of scenario s, denoted by $d(s)$, in place of the dose distribution $\tilde{d}(s)$ estimated by the static dose cloud approximation. In cases where the static dose cloud approximation holds, this means that our formulation will be equivalent to using a PTV; in cases where the static dose cloud approximation does not hold, our formulation provides a novel scenario-based robust planning method.

2.4. Beam-specific effects

The static dose cloud approximation does not model that the effect of an error can vary with the beams. Such variations can, however, be incorporated in the calculation of the weights $p_{i,s}$ by the definition of weights $p_{i,s}^{(b)}$ for each beam b in the set $\mathcal{B}$ of all beams, for all voxels i and scenarios s. The weights $p_{i,s}^{(b)}$ are under this refinement defined according to (2.4), but with $j^{(b)}$ substituted for j, where $j^{(b)}(s)$ is the offset in the dose grid corresponding to the shift of scenario s projected onto the beam plane of beam b. Thus, $j^{(b)}(s)$ neglects setup errors that are parallel to the beam axis, as illustrated in Figure 3. The weights for different beams are combined into a single weight $p_{i,s}$ by interpolation that takes into account the dose contribution to voxel i under scenario s from each beam. Let $d_i^{(b)}(s)$ denote the dose of beam b delivered to voxel i under scenario s. Then, the beam-specific counterpart of the criterion (2.2) takes the form

$$\sum_{s\in\mathcal{S}} \sum_{i\in\mathcal{C}} \phi(d_i(s)) \sum_{b\in\mathcal{B}} \frac{d_i^{(b)}(s)}{d_i(s)} p_{i,s}^{(b)}, \tag{2.5}$$

where $d_i^{(b)}(s)/d_i(s)$ is taken as $1/|\mathcal{B}|$ whenever $d_i(s) = 0$.

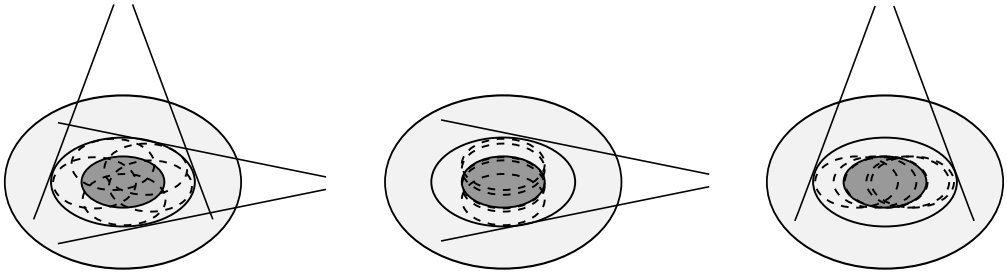

(a) $\mathcal{C}(s)$ under the scenarios s in $\mathcal{S}$.

(b) $\mathcal{C}(s)$ projected onto the beam plane for the beam along the x-axis.

(c) $\mathcal{C}(s)$ projected onto the beam plane for the beam along the y-axis.

Figure 3. Projection of $\mathcal{C}(s)$ (dashed ellipses) onto the beam planes. The projected sets are used in the calculation of beam-specific weights, but ϕ is still evaluated using the scenario doses without projections.

2.5. Extensions to other types of criteria

The reformulation (and consequently generalization, when d is substituted for $\tilde{d}$) of margins can be extended to other types of criteria than additively voxel-separable functions, as the following examples show: For a criterion that is a function ψ of an additively voxel-separable function, such as equivalent uniform dose measures [14], application of ψ to the left- and right-hand side of (2.3) yields the following reformulation:

$$\psi\left(\sum_{i\in\mathcal{P}}\phi(d_i)\right) = \psi\left(\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{C}} p_{i,s}\phi(\tilde{d}_i(s))\right).$$

Another common form of criteria are multiplicatively voxel-separable functions. Criteria of this type include tumor control probabilities, because the overall probability of tumor control is the product of the probability for tumor control in each tumor voxel. Derivations analogous to those in Section 2.2 yield the following reformulation for multiplicatively voxel-separable functions:

$$\prod_{i\in\mathcal{P}}\phi(d_i) = \prod_{s\in\mathcal{S}}\prod_{i\in\mathcal{C}}\phi(\tilde{d}_i(s))^{p_{i,s}}. \tag{2.6}$$

For a criterion of the form of a function ψ of a multiplicatively voxel-separable function, such as a normal tissue complication probability calculated using the relative seriality model [11], a reformulation is possible by application of ψ to the left- and right-hand side of (2.6). Finally, a constraint for all voxels i in $\mathcal{P}$ is equivalent to a constraint for all voxels i in $\mathcal{C}(s)$ and scenarios s in $\mathcal{S}$.

2.6. Experimental study

2.6.1. Optimization formulations and phantom geometries

We assessed the scenario-based generalization of margins by comparison against geometric margins and against previously published scenario-based robust planning

methods. The scenario-based methods that we used for benchmark were (composite) worst case optimization and expected value optimization [6, Sections II.D.1–2]. The comparison was performed with respect to stylized phantom geometries in either one dimension (1D) or two dimensions (2D), which contained a CTV with voxels enumerated by the set $\mathcal{C}$ and, for a subset of the cases, an OAR with voxels enumerated by the set $\mathcal{O}$. The geometries were planned for treatment with either a single orthogonal proton beam (incident from a dimension orthogonal to the patient geometry) or two coplanar proton beams at a perpendicular angle to each other. Treatment plan optimization using the various methods to achieve robustness was performed with respect to an objective function f penalizing deviation from a uniform dose of one to the CTV and a dose of zero elsewhere, according to

$$f(d) = w_{\mathcal{C}} \sum_{i \in \mathcal{C}} v_{i,\mathcal{C}} (d_i - 1)^2 + w_{\mathcal{O}} \sum_{i \in \mathcal{O}} v_{i,\mathcal{O}} d_i^2 + w_{\mathcal{E}} \sum_{i \in \mathcal{E}} v_{i,\mathcal{E}} d_i^2.$$

Here, $\mathcal{E}$ enumerates all voxels inside the external contour while for a general structure $\mathcal{R}$, $v_{i,\mathcal{R}}$ is the fraction of the structure that is contained in voxel i (this weight is used to take into account that the ROIs do not necessarily line up with the dose grid), and $w_{\mathcal{R}}$ is a scalar importance weight. Isotropic setup uncertainty of 1 cm was taken into account in the optimization. A total of four combinations of geometry, beam arrangement, and objective function were considered, which we labeled Examples 1–4, see Table 1.

The examples were selected to clearly exhibit the differences between the methods, and the methods are not necessarily intended to be used for such formulations: For example, the worst case optimization method is not intended to be applied to hopelessly conflicting scenarios, as in Examples 3 and 4. In such cases, a different set of scenarios should ideally be selected for the method to yield the desired result. Moreover, as explained for worst case optimization methods in a previous study [7], the methods are identical when there are no conflicts.

Table 1. Beam arrangement, structure definition, and objective function weights for the example cases. All intervals are specified in centimeters.

Example	Beam(s)	$\mathcal{C}$	$\mathcal{O}$	$\mathcal{E}$	$w_{\mathcal{C}}$	$w_{\mathcal{O}}$	$w_{\mathcal{E}}$
1	Orthogonal	$[-2, 2]$	–	$[-6, 6]$	10	–	1
2	Coplanar	$\left[-\frac{3}{2}, \frac{3}{2}\right]^2$	–	$[-6, 6]^2$	100	–	1
3	Orthogonal	$\left[-\frac{3}{2}, \frac{3}{2}\right]^2$	$[-5, -3] \times [-2, 2]$	$[-6, 6]^2$	100	100	1
4	Orthogonal	$\left[-\frac{3}{2}, \frac{3}{2}\right]^2$	$[-4, -2] \times [-2, 2]$	$[-6, 6]^2$	100	20	1

2.6.2. Numerical optimization and dose calculation

Proton beams were represented by scanning spots selected such that a Bragg peak position matched the voxel center for each voxel in the dose grid within 2 cm of the CTV. Spot doses for orthogonal beams were calculated using a Gaussian point-spread function with a standard deviation of 3 mm. Spot doses for coplanar beams

were calculated using the analytic proton depth-dose curve in Bortfeld [2]. The dose calculations were performed in a grid with a voxel side length of 1 mm. The set of possible setup shifts $\mathcal{S}$ was represented by 19 uniformly distributed scenarios in 1D (Example 1) and by nine scenarios (the nominal scenario, all shifts of 1 cm in the positive and negative axis directions, and all pairwise combinations of the axis directions) in 2D (Examples 2–4). A shift was modeled as a rigid translation of the beam isocenters relative to the patient coordinate system. The spot weights were optimized subject to nonnegativity bounds using the nonlinear programming solver SNOPT v7.2 (Stanford Business Software, Palo Alto, California).

Treatment plan optimization using scenario-based margins was implemented on the beam-specific form of Section 2.4 for the beam arrangement with two fields (Example 2). Optimization using geometric margins was implemented as minimization of the objective function value in the nominal scenario evaluated with respect geometrically expanded structures. The CTV was expanded to a PTV and the OAR expanded to a PRV by substitution of an expanded volume defined as $\{i + j(s) : i \in \mathcal{R}, s \in \mathcal{S}\}$ for the original structure $\mathcal{R}$. The objective function weights were scaled with a factor equal to the fractional increase in volume due to the expansion in order to obtain comparable results with the scenario-based methods. The maximum operator used to compute the worst case objective function value in the worst case optimization was approximated by a smooth power mean according to

$$\max_{s \in S} f(d(s)) \approx \left(\frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} f(d(s))^{10} \right)^{1/10}$$

in order for the operator to be continuously differentiable. The expected objective function value used in the expected value optimization was calculated with respect to normally distributed errors with standard deviation selected such that 95% of all values were contained within 1 cm, conditioned on that an error of at most 1 cm would occur.

3. Results

Doses are quantified in percent of the prescription. We use the term "margins" to refer both geometric and scenario-based margins when their results are similar.

3.1. Example 1

This example compares the previous scenario-based methods against margins under very simple conditions: the static dose cloud approximation is exact and there is no OAR that competes with target coverage. The results for geometric and scenario-based margins coincided precisely.

Figure 4 shows that all methods gave a uniform dose of about 97% to the innermost 2 cm of the CTV. Margins kept the planned dose uniform at this level inside the entire PTV, except for at the very edges where the dose was increased to form "shoulders." Worst case optimization created a mildly decreasing dose with increasing distance to the origin outside the innermost 2 cm of the CTV. Like margins, it

created shoulders at the edges of the high-dose region. Expected value minimization created a smooth and gradual fall-off to a dose of about 30% at the edges of the PTV, without any shoulder formations.

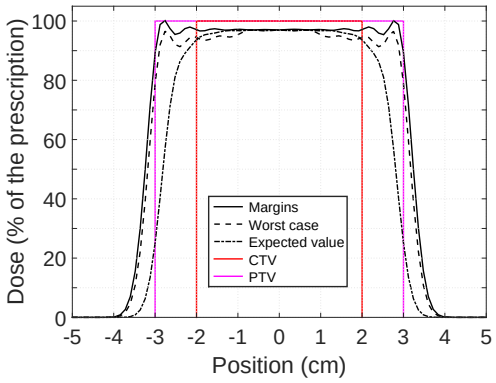

Figure 4. Nominal doses for Example 1. The height of the CTV and PTV indicates the prescription dose.

3.2. Example 2

This example highlights the differences between the scenario-based methods and geometric margins if the static dose cloud approximation is inaccurate. The reason that the static dose cloud approximation gives a poor prediction of a perturbed dose is that the beam doses are unaffected by components of the shifts that are parallel to the beam axis (neglecting the small effect that the beam doses are scaled according to the inverse-square law).

Figure 5 depicts the nominal dose for scenario-based and geometric margins. Associated DVHs for all scenarios are shown in Figure 6. Corresponding results are not shown for worst case and expected value optimization because of the close similarity with scenario-based margins. However, line doses for all methods are depicted in Figure 7. Geometric margins covered the PTV with a uniform dose at the prescription level. The scenario-based methods instead created a linearly decreasing dose from 100% at the CTV edge to 90% at the proximal edge of the PTV and to about 70% at the distal edge of the PTV. The CTV dose of the scenario-based methods was despite the lower integral dose uniform around 100% in all scenarios, whereas the CTV dose varied in the range 90–135% for geometric margins. The beam doses in Figure 8 explain why the plan created by geometric margins is sensitive to perturbations: The saddle-shaped doses (peaks on the lateral edges of the CTV and valleys on the proximal and distal edges) sum to a highly uniform total dose with a steep fall-off in the nominal scenario, but the total dose is irregular if the beams do not patch as intended. The scenario-based methods created uniform beam doses that do not rely on that the beams patch accurately.

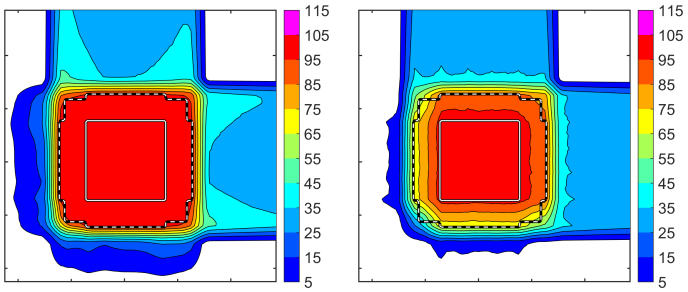

(a) Geometric margins (b) Scenario-based margins

Figure 5. Nominal doses for Example 2. The CTV is indicated by a solid line and the PTV is indicated by a dashed line.

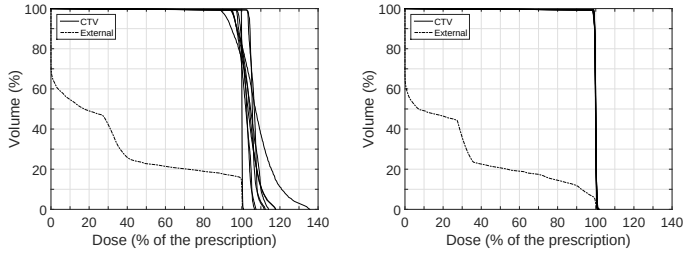

(a) Geometric margins (b) Scenario-based margins

Figure 6. DVHs for Example 2 in all scenarios.

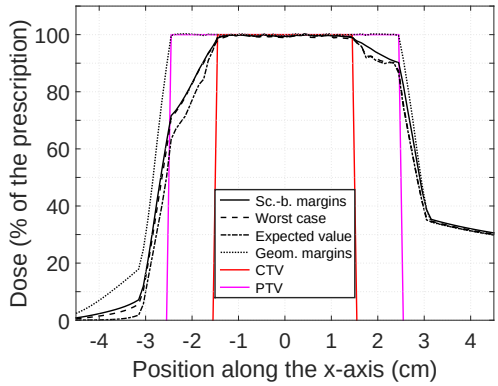

Figure 7. Line dose profiles for Example 2 in the nominal scenario for a trace along the x-axis through the isocenter.

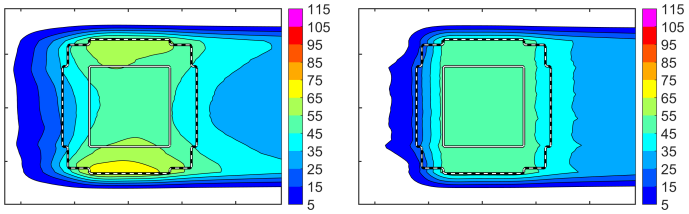

(a) Geometric margins (b) Scenario-based margins

Figure 8. Beam doses for Example 2 for the beam along the x-axis. The beam doses for the beam along the y-axis were identical except for a 90° rotation. The CTV is indicated by a solid line and the PTV is indicated by a dashed line.

3.3. Example 3

The purpose of this example is to investigate how scenario-based margins handle a situation when target coverage in some scenarios is incompatible with OAR sparing. The incompatibility stems from that the PTV overlaps with the PRV. The static dose cloud approximation is close to exact: it would have been exact had it not been for shift scenarios in pairwise combinations of the axis direction that do not line up exactly with the dose grid. The results for geometric margins were, nevertheless, highly similar to the results for scenario-based margins (98% of the differences in voxel dose were less than 3% of the prescription) and are therefore not shown.

Figure 9 shows that margins and expected value optimization extended the high-dose region from the CTV to the edges of the PTV except for in the leftward direction that is in conflict with the OAR. The CTV coverage was therefore maintained under all shifts without rightward component. Coverage was lost almost completely for 20% of the CTV under rightward shifts, see Figure 10. A minor difference for expected value optimization compared to margins was that the doses in the CTV-to-PTV margin were slightly lower and the dose fall-off was more gradual. Worst case optimization only extended the high-dose region to the right of the CTV. The CTV coverage was therefore poor not only under the rightward shifts that are in conflict with OAR sparing, but also under the relatively "easier" up- and downward shifts. The line doses in Figure 11 demonstrate how the high-dose region for worst case optimization was retracted in the up- and downward directions compared to the other methods.

3.4. Example 4

The difference between this example and Example 3 is that the PRV overlaps with the CTV and the weight on the OAR objective is five times smaller than previously. The results for geometric margins are not depicted because of close similarity with the results for scenario-based margins (98% of the differences in voxel dose were less than 3% of the prescription).

The lesser priority on OAR sparing compared to Example 3 made margins accept a dose of about 80% in the overlap region between the PTV and PRV, see Figure 12. The dose compromise in this region was controlled in the sense that the dose was

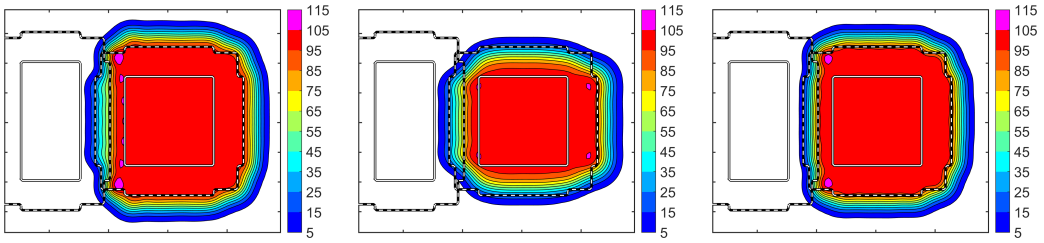

(a) Scenario-based margins (b) Worst case optimization (c) Expected value optimization

Figure 9. Nominal doses for Example 3. The CTV and OAR are indicated by solid lines and the PTV and PRV are indicated by dashed lines.

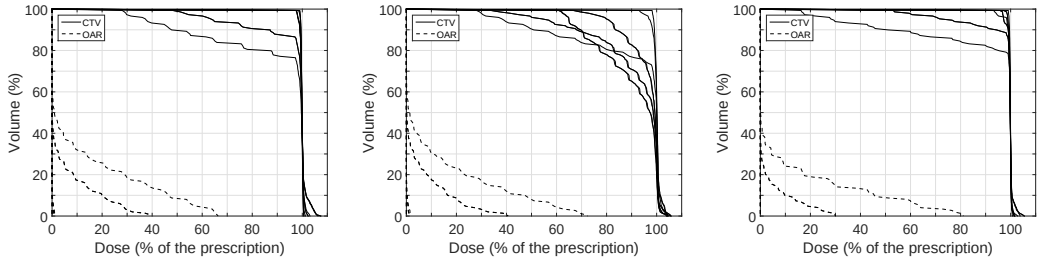

(a) Scenario-based margins (b) Worst case optimization (c) Expected value optimization

Figure 10. DVHs for Example 3 in all scenarios.

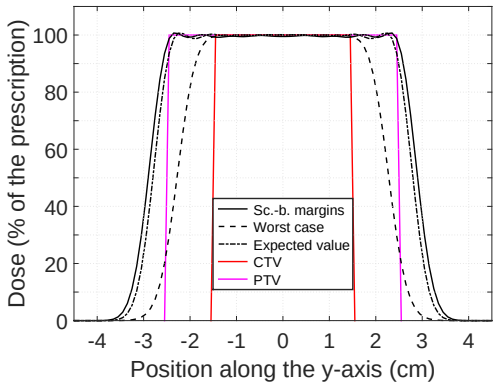

Figure 11. Line dose profiles for Example 3 in the nominal scenario for a trace along the y-axis through the isocenter. The height of the CTV and PTV indicates the prescription dose.

uniform, except for peaks in the dose at the left edge of the PTV. Worst case optimization gave a uniform dose at the prescription level to the CTV in the left shift scenario, and underdosed the CTV to about 70% in the region where it overlaps with the PRV. The CTV coverage was poor also under all shifts with an up- or downward component because there was little extension of the high-dose region outside the CTV in these directions, see Figure 13. Expected value optimization covered almost the entire CTV with a dose of one in the nominal scenario, thereby accepting doses above 90% to more than 20% of the OAR if a leftward shift occurs. The CTV dose was highly non-robust to shifts in the right direction because the dose falls gradually towards zero left of the CTV edge, see Figure 14. The CTV coverage under a shift with leftward component was comparable to margins.

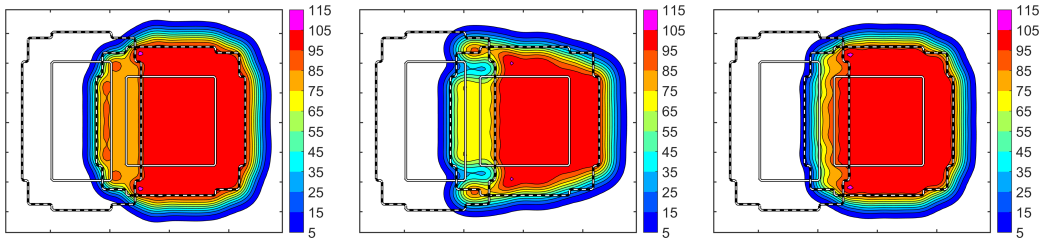

(a) Scenario-based margins (b) Worst case optimization (c) Expected value optimization

Figure 12. Nominal dose distributions for Example 4. The CTV and OAR are indicated by solid lines and the PTV and PRV are indicated by dashed lines..

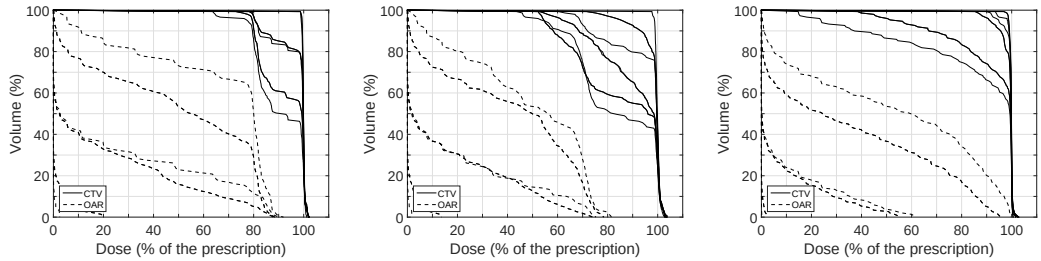

(a) Scenario-based margins (b) Worst case optimization (c) Expected value optimization

Figure 13. DVHs for Example 4 in all scenarios.

4. Discussion

4.1. Advantages over geometric margins

There is clear evidence that scenario-based optimization can provide better target coverage robustness than margins and also reduce the dose to healthy tissues if the static dose cloud approximation is inaccurate [3,6,8,12]. Scenario-based optimization is therefore particularly advantageous for ion therapy, treatments of heterogeneous densities, and treatments with few fields. Our results for Example 2 reinforce this conclusion: geometric margins created nonuniform beam doses that were sensitive

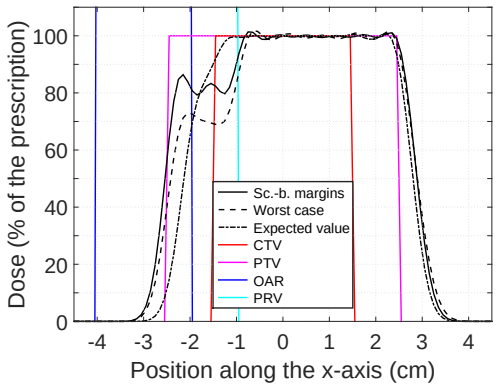

Figure 14. Line dose profiles for Example 4 in the nominal scenario for a trace along the x-axis through the isocenter. The height of the CTV and PTV indicates the prescription dose. The OAR and PRV boundaries are indicated by vertical lines.

to perturbations and placed more dose inside the CTV-to-PTV margin than necessary, while all scenario-based methods created highly robust plans. Furthermore, a potential advantage of scenario-based methods, even if scenario doses are calculated using the static dose cloud approximation, is that robustness can be implemented for nonuniform CTV prescriptions, as in dose-painting guided by functional imaging [1].

4.2. Advantages over other scenario-based models

4.2.1. Dose gradients

This study, alongside previous data [6], shows that expected value optimization with respect to systematic setup uncertainty can lead to doses that fall gradually from the prescription level to a dose of zero over a prolonged distance outside the CTV. Such a gradual fall-off was observed consistently in Examples 1, 3, and 4. Margins and worst case optimization instead created dose shoulders at the PTV boundary that contributed to a steep fall-off. The gentle fall-off created by expected value optimization is a point of concern in view of the fact that the possibility to deliver dose-escalated treatments without compromising the dose to surrounding organs has been a key driver in the adoption of intensity-modulated over conformal radiation therapy [17]: van Herk et al. [22] showed that steep dose gradients are necessary for such dose escalation. Proponents of proton therapy also often point to the steep fall-off on the lateral and distal edge of the target as evidence of its superiority over photon treatments [18].

4.3. Tradeoffs between conflicting criteria

In a previous study [7], we showed that worst case optimization can be too conservative if target coverage to some extent is incompatible with sparing of an OAR. The optimization can then neglect coverage also in scenarios that are not in direct conflict with the OAR. This effect was also apparent in Examples 3 and 4, where worst case optimization disregarded CTV coverage with respect to shifts with an up-

or downward component. Margins and expected value optimization did not exhibit this shortcoming. These methods instead created robust CTV coverage for all shifts except those that required an extension of the high-dose region into the PRV. For the worst case optimization method to work as intended, the scenarios of non-resolvable conflict could be retracted.

Margins and expected value optimization exhibited markedly different behaviors with respect to the dose compromise in the overlap region between the PTV and PRV in Example 4. Margins created a controlled underdosage of the CTV while expected value optimization sacrificed CTV coverage under a rightward shift and OAR sparing under a leftward shift. This is an effect of that the nominal scenario has a larger probability of occurring than the shift scenarios. Furthermore, the results for expected value optimization depend on the scenario discretization: The probability for the extreme shift scenarios become smaller if intermediate shifts are also incorporated in the uncertainty set $\mathcal{S}$ due to normalization of the discretized probabilities.

It is a clinical decision whether the characteristics of margins or expected value optimization are preferable. The current ICRU guidelines advocate that when the PTV and OAR or PRV overlap, priority rules should be used, or the regions should be subdivided into regions with different absorbed-dose constraints [10, Section 4.5]. This effect is achieved by margin-based planning with importance weights for the PTV and the OAR or PRV used in the optimization. Several clinical protocols, however, prescribe a smaller margin in the interface against some radiation sensitive structures, for example a retracted CTV-to-PTV margin for prostate in the posterior direction to promote rectal sparing [5, 24]. It is our view that if a retraction is to be made in a scenario-based model, it should be incorporated in the selection of scenarios rather than in their probabilities, in order to maintain steep dose gradients. Optimization of the definition of the uncertainty set $\mathcal{S}$ subject to constraints that a set of clinical goals is satisfied for all scenarios in $\mathcal{S}$ is described in Fredriksson et al. [9].

4.4. Potential disadvantages

4.4.1. Nonconvexity

The scenario-based margins lead to a nonconvex optimization model if beam-specific effects are taken into account, as shown in Appendix A. Nonconvex optimization problem can have multiple local minima, meaning that optimization algorithms cannot guarantee convergence to the global optimum. While a nonconvex formulation allows the existence of multiple local optima, however, it does not assure that such optima exist. The local optima may also be of negligible clinical difference even if they do exist.

4.4.2. Limitation to setup errors in homogeneous media

The beam-specific form of the scenario-based margins does not take the mass density of the treatment volume into account in the calculation of the beam-, scenario- and

voxel-specific weights. The model is therefore appropriate for treatments with few fields, as demonstrated by Example 2, but although the correct scenario doses are used, the weights may be inappropriate for irradiation of heterogeneous media. The scenario-based margins as presented here are also not designed for handling of ion beam range errors, and cannot utilize multiple patient images for handling of inter- and intra-fraction organ motion. Extensions to these types of uncertainty will be presented in a separate publication.

5. Conclusions

We have introduced a scenario-based generalization of treatment plan optimization with margins to handle setup uncertainties. When the static dose cloud approximation is exact, the generalization is equivalent to margins, but when the approximation is inaccurate, the generalization exhibits the benefits of scenario-based optimization. Comparative planning on phantom geometries demonstrated that the scenario-based margins can provide robustness where geometric margins fail and at the same time avoid some disadvantages of previous methods for robust planning, such as the unnecessarily gentle dose fall-off of expected value optimization and the neglection of "easy" scenarios of worst case optimization.

A. Proof of nonconvexity

The following example shows that a beam-weighted criterion f defined according to (2.5) can be nonconvex even when ϕ is convex. Consider a 2D geometry where a single voxel is irradiated by a first beam along the x-axis and a second beam along the y-axis. Let $\phi(d) = |d|^{3/2}$ and let $\mathcal{S}$ be an uncertainty set composed of the nominal scenario and a 1-voxel shift along the y-axis. The beam- and scenario-specific weights then take on the values $p_s^{(1)} = 1$ and $p_s^{(2)} = 1/2$ for all s in $\mathcal{S}$. Consider two plans A and B with beam doses according to

$$d_{\mathrm{A}}^{(b)}(s) = \begin{cases} 2 & \text{if } b = 1 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad d_{\mathrm{B}}^{(b)}(s) = \begin{cases} 4 & \text{if } b = 2 \\ 0 & \text{otherwise} \end{cases}, \quad \forall s \in \mathcal{S}.$$

Convexity of f requires that the function value of the average dose is less than or equal to the average function value. The inequality

$$f((d_{\mathrm{A}} + d_{\mathrm{B}})/2) = 4\sqrt{3} > (f(d_{\mathrm{A}}) + f(d_{\mathrm{B}}))/2 = 2(\sqrt{2} + 2),$$

therefore, proves the nonconvexity of f.

References

[1] S. Bentzen. Theragnostic imaging for radiation oncology: dose-painting by numbers. *Lancet Oncol.*, 6(2):112–117, 2005.

[2] T. Bortfeld. An analytical approximation of the Bragg curve for therapeutic proton beams. *Med. Phys.*, 24(12):2024–2033, 1997.

[3] W. Chen, J. Unkelbach, A. Trofimov, T. Madden, H. Kooy, T. Bortfeld, and D. Craft. Including robustness in multi-criteria optimization for intensity-modulated proton therapy. *Phys. Med. Biol.*, 57(3):591–608, 2012.

[4] T. Craig, J. Battista, and J. Van Dyk. Limitations of a convolution method for modeling geometric uncertainties in radiation therapy. I. The effect of shift invariance. *Med. Phys.*, 30(8):2001–2011, 2003.

[5] D. Dearnaley, I. Syndikus, G. Sumo, M. Bidmead, D. Bloomfield, C. Clark, A. Gao, S. Hassan, A. Horwich, R. Huddart, V. Khoo, P. Kirkbride, H. Mayles, P. Mayles, O. Naismith, C. Parker, H. Patterson, M. Russell, C. Scrase, C. South, J. Staffurth, and E. Hall. Conventional versus hypofractionated high-dose intensity-modulated radiotherapy for prostate cancer: preliminary safety results from the CHHiP randomised controlled trial. *Lancet Oncol.*, 13(1):43–54, 2012.

[6] A. Fredriksson. A characterization of robust radiation therapy treatment planning methods—from expected value to worst case optimization. *Med. Phys.*, 39(8):5169–5181, 2012.

[7] A. Fredriksson and R. Bokrantz. A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning. *Med. Phys.*, 41(8):081701, 2014.

[8] A. Fredriksson, A. Forsgren, and B. Hårdemark. Minimax optimization for handling range and setup uncertainties in proton therapy. *Med. Phys.*, 38(3):1672–1684, 2011.

[9] A. Fredriksson, A. Forsgren, and B. Hårdemark. Maximizing the probability of satisfying the clinical goals in radiation therapy treatment planning under setup uncertainty. *Med. Phys.*, 42(7):3992–3999, 2015.

[10] International Commission on Radiation Units and Measurements. Report 83: Prescribing, recording, and reporting photon-beam intensity-modulated radiation therapy (IMRT). *J. ICRU*, 11(1), 2010.

[11] P. Källman, A. Ågren, and A. Brahme. Tumour and normal tissue responses to fractionated non-uniform dose delivery. *Int. J. Radiat. Biol.*, 62(2):249–262, 1992.

[12] W. Liu, X. Zhang, Y. Li, and R. Mohan. Robust optimization of intensity modulated proton therapy. *Med. Phys.*, 39(2):1079–1091, 2012.

[13] A. Lomax. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 2: the potential effects of inter-fraction and inter-field motions. *Phys. Med. Biol.*, 53(4):1043–1056, 2008.

[14] A. Niemierko. A generalized concept of equivalent uniform dose (EUD). *Med. Phys.*, 26(6):1100, 1999.

[15] U. Oelfke and T. Bortfeld. Inverse planning for photon and proton beams. *Med. Dosim.*, 26(2):113–124, 2001.

[16] D. Pflugfelder, J. Wilkens, and U. Oelfke. Worst case optimization: a method to account for uncertainties in the optimization of intensity modulated proton therapy. *Phys. Med. Biol.*, 53(6):1689–1700, 2008.

[17] N. Sheets, G. Goldin, A.-M. Meyer, Y. Wu, Y. Chang, T. Stürmer, J. Holmes, B. Reeve, P. Godley, W. Carpenter, and R. Chen. Intensity-modulated radiation therapy, proton therapy, or conformal radiation therapy and morbidity and disease control in localized prostate cancer. *JAMA*, 307(15):1611–1620, 2012.

[18] H. Suit, H. Kooy, A. Trofimov, J. Farr, J. Munzenrider, T. DeLaney, J. Loeffler, B. Clasie, S. Safai, and H. Paganetti. Should positive phase III clinical trial data be required before proton beam therapy is more widely adopted? No. *Radiother. Oncol.*, 86(2):148–153, 2008.

[19] J. Unkelbach, T. Bortfeld, B. Martin, and M. Soukup. Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning. *Med. Phys.*, 36(1):149–163, 2009.

[20] J. Unkelbach, T. Chan, and T. Bortfeld. Accounting for range uncertainties in the optimization of intensity modulated proton therapy. *Phys. Med. Biol.*, 52(10):2755–2773, 2007.

[21] M. van Herk. Errors and margins in radiotherapy. *Semin. Radiat. Oncol.*, 14(1):52–64, 2004.

[22] M. van Herk, P. Remeijer, and J. Lebesque. Inclusion of geometric uncertainties in treatment plan evaluation. *Int. J. Radiat. Oncol.*, 52(5):1407–1422, 2002.

[23] M. van Herk, P. Remeijer, C. Rasch, and J. Lebesque. The probability of correct target dosage: dose-population histograms for deriving treatment margins in radiotherapy. *Int. J. Radiat. Oncol.*, 47(4):1121–1135, 2000.

[24] M. Zelefsky, D. Crean, G. Mageras, O. Lyass, L. Happersett, C. Ling, S. Leibel, Z. Fuks, S. Bull, H. Kooy, M. van Herk, and G. Kutcher. Quantification and predictors of prostate position variability in 50 patients evaluated with multiple CT scans during conformal radiotherapy. *Radiother. Oncol.*, 50(2):225–234, 1999.